

REPORT DOCUMENTATION PAGE

a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED		1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution is unlimited	
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE			
4. PERFORMING ORGANIZATION REPORT NUMBER(S)		5. MONITORING ORGANIZATION REPORT NUMBER(S)	
6a. NAME OF PERFORMING ORGANIZATION Computer Science Dept. Naval Postgraduate School		6b. OFFICE SYMBOL (if applicable) CS	
7a. NAME OF MONITORING ORGANIZATION Naval Postgraduate School			
c. ADDRESS (City, State, and ZIP Code) Monterey, CA 93943-5000		7b. ADDRESS (City, State, and ZIP Code) Monterey, CA 93943-5000	
8a. NAME OF FUNDING/SPONSORING ORGANIZATION		8b. OFFICE SYMBOL (if applicable)	
9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER			
c. ADDRESS (City, State, and ZIP Code)		10. SOURCE OF FUNDING NUMBERS	
		PROGRAM ELEMENT NO.	PROJECT NO.
		TASK NO.	WORK UNIT ACCESSION NO.
1. TITLE (Include Security Classification) HYPER-NPSNET: A VIRTUAL WORLD WITH AN INTEGRATED 3D HYPERTEXT (U)			
2. PERSONAL AUTHOR(S) Daley, John A.			
3a. TYPE OF REPORT Master's Thesis	13b. TIME COVERED FROM 10/89 TO 03/92	14. DATE OF REPORT (Year, Month, Day) 26 March 1992	15. PAGE COUNT 53
5. SUPPLEMENTARY NOTATION The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the United States Government.			
7. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	
		Virtual world, hypertext, hypermedia, virtual reality, artificial reality, 3D, cyberspace	
9. ABSTRACT (Continue on reverse if necessary and identify by block number) This thesis proposes an extension to the NPSNET 3D virtual world prototype to provide an integrated 3D hypertext. This hypertext would be embedded into the virtual world, and would provide the capability for real-time or non-real-time reference, aspect-change within or around vehicles, direct navigation to a specified 3D location, and other features. Hypermedia such as video or graphic animations, with or without sound, in addition to text files will be accessible through the hypertext. The hypertext elements will be embedded into the virtual world by the use of icons or other symbols, which become visible if selected by the user. Foreseeable applications include hypernavigation, historical reference visualization, direct access to higher-level analysis functions, and contextual referencing for training applications. A prototype of Hyper-NPSNET is planned for initial review in June 1992.			
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS		21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED	
22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. Michael Yzda		22b. TELEPHONE (Include Area Code) (408) 646-2305	22c. OFFICE SYMBOL CS/Zk

T260065

Approved for public release; distribution is unlimited

**HYPER-NPSNET:
A VIRTUAL WORLD WITH AN INTEGRATED 3D HYPERTEXT**

by

John Alexander Daley
Lieutenant Commander, United States Navy
B.S., University of Florida, 1979

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF COMPUTER SCIENCE

from the

NAVAL POSTGRADUATE SCHOOL
26 March 1992

ABSTRACT

This thesis proposes an extension to the NPSNET 3D virtual world prototype to provide an integrated 3D hypertext. This hypertext would be embedded into the virtual world, and would provide the capability for real-time or non-real-time reference, aspect-change within or around vehicles, direct navigation to a specified 3D location, and other features. Hypermedia such as video or graphic animations, with or without sound, in addition to text files will be accessible through the hypertext. The hypertext elements will be embedded into the virtual world by the use of icons or other symbols, which become visible if selected by the user. Foreseeable applications include hypernavigation, historical reference visualization, direct access to higher-level analysis functions, and contextual referencing for training applications. A prototype of Hyper-NPSNET is planned for initial review in June 1992.

TABLE OF CONTENTS

I.	THE NEED FOR HYPERTEXT	1
A.	APPLICATIONS AND CHARACTERISTICS OF VIRTUAL WORLDS	1
B.	VIRTUAL WORLD INTERFACE DEVICES	3
C.	NPSNET	3
D.	HYPER-NPSNET	5
II.	PREVIOUS WORK	7
A.	HYPERTEXT AND HYPERMEDIA DEFINED	7
B.	VIRTUAL WORLDS	7
C.	EXAMPLES OF VIRTUAL WORLDS WITH HYPERTEXT	8
1.	The Information Visualizer	9
D.	ADVANTAGES OF COMBINING VIRTUAL WORLDS AND HYPERTEXT	10
III.	3D/4D HYPERTEXT	12
A.	3D/4D HYPERTEXT DEFINED	12
B.	LOCATION OF HYPERTEXT NODES IN 3-SPACE	13
C.	SELECTION OF 3D HYPERTEXT NODES	14
D.	HYPERSEARCH	16
1.	3D Zones	17
a.	Types of 3D Zones	17
b.	Type of Applications Utilizing 3D Zones	17
c.	Flexible Hypertext Links	18
d.	An Application of 3D Zones	18
2.	An Application of HyperSearch	19
E.	HYPERNAVIGATION	20
1.	Hypernavigation Defined	20
2.	Portals	20
3.	The “Alice” Effect	21
4.	Hypernavigation Interface	22
5.	Three-dimensional traversal tracking	22
F.	INSTANT ASPECT CHANGE	23

IV. AUTHORIZING IN HYPER-NPSNET	25
A. THE NEED FOR AGILE AUTHORIZING	25
B. OBJECT-ORIENTED CONSTRUCTION	25
C. USER-LEVEL AUTHORIZING	26
D. DESIGN-LEVEL AUTHORIZING	27
E. AN EXAMPLE OF AGILE AUTHORIZING	28
V. APPLICATIONS OF HYPER-NPSNET	29
A. CHARACTERISTICS OF SUITABLE APPLICATIONS	29
1. Fighting the Virtual Battle.....	29
2. Magnetic Disk File Maintenance Training.....	31
3. DSRV Piloting.....	32
VI. SUMMARY AND CONCLUSIONS	34
A. VIRTUAL WORLD DEVELOPMENT	34
B. MODERN HYPERTEXT TRENDS	35
C. VIRTUAL WORLD INTERFACE ISSUES	35
D. FOUR-DIMENSIONAL HYPERTEXT	36
REFERENCES	38
ANNOTATED BIBLIOGRAPHY	42
INITIAL DISTRIBUTION LIST	45

ACKNOWLEDGEMENTS

Many have contributed to the formulation of this thesis, and I would like to sincerely thank all that have helped in my efforts, whether their contribution was that magical bit of inspiration that leads to a blinding flash of the obvious, or a cold slap in the face to “inspire” me to get it finished. I am grateful to you all.

I am particularly indebted to Dr. Michael Zyda, whose belief in me, first and foremost, gave me the confidence to tackle a project that had almost no historical precedent. Mike helped me through the tough times, and there were many, when writing at length about the hypothetical. His pointing to the right sources and references saved me a great deal of time, and proved invaluable in my research.

Dr. Yuh-jeng Lee tapped a spring of interest within me in the area of Intelligent Tutoring Systems (ITS's) and education in general. His fine teaching and his counsel helped me find a deep passion within me for improving the lot of those being educated on the planet. I hope we succeed.

To all in the Naval Postgraduate School Macintosh User Group, and particularly to Frame Technologies and their fine product, FrameMaker for the Macintosh, I extend my admiration and my thanks.

Finally to my wife Margie and my four sons Vincent, Kyle, Matthew, and Charles, I appreciate deeply the patience you have shown while I have been a stranger to you, and this following a submarine tour at sea. Your support while I battled thesisitis made it possible for me to succeed.

I. THE NEED FOR HYPERTEXT

Virtual world programs provide the user of a computer system the capability of real-time interaction with objects in a simulated environment. This virtual world can either be a simulation of an actual environment (virtual reality), or an artificially-created environment for specialized applications (cyberspace). An essential feature of virtual world applications is the ability of the user to interact with the created environment in a non-scripted, real-time, interactive manner.

A. APPLICATIONS AND CHARACTERISTICS OF VIRTUAL WORLDS

An extremely wide range of diverse applications exist for virtual world programs. Since virtual world technology is still in its infancy and has not yet been embraced by the commercial market at large, many applications that could effectively employ virtual world technology have yet to be developed. The computational limitations and expense of many current workstations preclude widespread marketing of virtual world applications, although some recent work demonstrates the ability to employ off-the-shelf hardware in virtual world applications and design [Zyda91, Zyda92]. With the recent advent of 3D rendering programs for personal computers [Gore92, McManus92] and low-cost 3D workstations [Stahr91], virtual world designers now have the capability to write applications for a burgeoning market. Some applications of virtual world systems are described below.

Intelligent tutoring systems using virtual worlds can simulate training scenarios and interact with the user to enhance performance skills. Taking advantage of realistic 3D representation of the desired training environment (especially in applications involving moving platforms, such as flight simulators or robotic devices), the trainee can realistically interact with objects in a virtual world without being constrained by restrictions that exist in the physical world. An engineer can be trained on a nuclear propulsion plant, for

example, that offers the capability to travel to the various engineering spaces at an instant in time to see the status of the plant from various perspectives. The same engineer could even enter the virtual reactor vessel and examine fluid levels, navigate the spaces looking for primary leaks, or go to the scene of a casualty to determine current status, all by traversing the virtual training environment. By remembering the training simulations completed by each student and recording any incorrect actions performed, the virtual intelligent tutoring system can tailor future training scenarios, or even select the appropriate virtual environment, to optimize the training benefit for each trainee [Woolf90]. This capability can exploit virtual world technology to perform real-time training, by remembering where the student was in the virtual world when errors were made or when training was interrupted.

Virtual world Computer-Aided Design (CAD) programs can model enhancements to existing products, or create new ones, that can be designed in cyberspace before actually existing. With relatively low-cost workstations emerging in the marketplace, some systems can be designed by the prospective users, instead of design teams that are not familiar with the environment where the product will be used. For example, submarine fire-control systems that perform presetting and control functions on torpedoes, cruise missiles, mines, and other weapons have long suffered interface problems with the actual crew members who use them, partly due to the lack of submarine experience by the design engineers [Daley85]. With virtual world modeling of the new fire control system's interface, the users can operate and modify the new system before it is ever built, saving costly and time-consuming alterations to the fire control system and the resultant training expense and performance degradation. Obtaining this type of efficiency by eliminating wasteful modification expense is, in fact, the goal of the United States Department of Defense Modeling and Simulation Plan [Atwood91].

3D graphics programs can use a virtual world to present data in a new and effective way, allowing the user to interact with the data itself in real time [Brand88, Card91]. A field commander can analyze updated intelligence information by actually visualizing reports of

enemy strength in a virtual environment; for example, tank and artillery strength. By modeling what were originally raw numbers into the actual vehicles, and placing them in the applicable environment, the commander can better visualize enemy strength, adjust battle tactics, and perform “what-if” analysis on the data, including manipulation of the time variable to analyze trends and project future engagements. Xerox PARC’s Information Visualizer provides the capability to compress or expand time while viewing chronological events [Mackinlay91, Card91, Henderson86]. Some of the extensions discussed below can add the ability for the user to directly access related information elements (files, graphics, sounds, animations) that are viewed on the chronological perspective wall.

B. VIRTUAL WORLD INTERFACE DEVICES

Many virtual world programs incorporate special interface devices to provide the user with the capability of effective 3D interaction with the simulated environment [Krueger91, Rheingold91]. These devices include stereo goggles, data gloves, and body suits. While these devices may be necessary for certain applications, many virtual world applications have no need of these currently expensive and cumbersome devices: the user can effectively interact in 3D using a 2D CRT that simulates 3D. The use of special 3D input devices, such as a spaceball or flying mouse, is not inherent to the definition of a virtual world; however, these devices can facilitate interaction with the environment in a more efficient and realistic manner. This thesis does not assume the use of these devices for virtual world applications unless so stated.

C. NPSNET

NPSNET is a virtual world system developed at the Naval Postgraduate School in Monterey, California [Marsan91, Zyda91, Zyda92]. In its current configuration, NPSNET runs on Silicon Graphics IRIS workstations and uses a spaceball as a 3D input device, in addition to the more traditional button, knob, and mouse inputs. The architecture of

NPSNET provides the capability for multiple users to simultaneously interact with the environment in real time over an Ethernet network.




Figure 1.1 NPSNET Interface Example

NPSNET (see figure above) models a 50 kilometer square geographic area located at Fort Hunter-Liggett, California. This area is modeled using actual detailed data from the terrain, including hills, valleys, and lakebeds. The program models a variety of vehicles both real and imaginary, including tanks, aircraft, and a variety of land vehicles. The user can operate any vehicle not currently being driven by another user, and can fire on any other vehicle (which can shoot back). Autonomous vehicles operate in the background, and can be interacted with by any user on the network.

NPSNET shows that a real-time, interactive 3D simulation system for multiple simultaneous participants is achievable using low-cost workstations. The user can enter one of 500 active vehicles and interact with other users and autonomous vehicles. Real-time

vehicle movement and texturing of objects is modeled in the virtual world. Vehicle positions are predicted using dead-reckoning until the user changes direction or speed, and updated positions of user-driven vehicles are passed over the Ethernet. Workstations used in NPSNET range from the relatively-inexpensive IRIS Indigo to an IRIS model VGX with multiple processors¹ [Zyda91].

Planned applications for NPSNET include training and strategy development for coordinated operations and war gaming. Other environments can be incorporated in the future, including ocean environments for naval surface and subsurface simulations.

D. HYPER-NPSNET

Hyper-NPSNET originated with a desire to embed information into the terrain, to provide for analysis that extended beyond interaction with the virtual objects. By providing buttons or other icons that appear, by selectable option, in the virtual environment, the user can retrieve information about fixed or moving objects, observe or effect changes in the environment, and perform high-level analysis on objects in the virtual world. The information obtained by the user can take the form of text, graphics, animation, sound, or a combination of the above, and can be linked to other related information. This information can be accessed either in real-time, compressed or expanded time, or in a time-frozen state, to provide look-up capability for users in a training scenario, for example.

While each information element could be individually embedded into the terrain, the resulting structure does not provide the ability to connect the embedded elements. Since the embedded elements could take the form of a reference manual, for example, that other elements require the use of, it became clear that some of these embedded elements would need to be accessed by more than one item in the virtual world. The need to link these information elements together suggests the need for a hypertext/hypermedia structure [Nielsen90]. Additionally, in very large systems, the requirement to access each item

1. The Personal IRIS model does not have texturing capability, but otherwise has the identical functionality on NPSNET as the more expensive workstations.

individually is tedious, and causes excessive icon clutter in the virtual world. By linking these information elements together and embedding each of them into a specific location in the virtual world, an underlying hypertext provides the ability to explore related items in the hypertext via an information retrieval system; provide search analysis capability for objects in the environment, adjust viewing perspective to predefined specific vantage points, and even directly travel to other locations in the virtual world. The resulting *HyperWorld* has the power of hypertext not alongside a virtual world, but embedded into it. Hyper-NPSNET is the planned implementation of an embedded hypertext into the current NPSNET virtual world.

II. PREVIOUS WORK

A. HYPERTEXT AND HYPERMEDIA DEFINED

As contrasted with normal text which is sequential, *hypertext* is non-sequential in its structure, with *links* connecting the non-sequential *nodes* of text [Nielsen90, pp. 1-3]. Rather than force the reader to consecutively proceed from sentence to sentence or page to page, with a hypertext the reader can be given multiple options after reading a given chunk of text. For example, after reading an introduction to calculus in a hypertext-driven training system, the reader can be given options to start the section on derivatives, consider present-day practical examples of the application of calculus, review pre-calculus material, or send electronic mail to the professor (possibly to express panic or concern). *Hypermedia* extends the concept of hypertext to include non-textual nodes, such as graphics, pictures, sound, animations, or combinations of these. Some of these nodes may actually be applications, such as a program that performs calculations on data that exists in a particular node at a given point in time.

B. VIRTUAL WORLDS

The mention of the term "virtual world" to some creates an image of 3D goggles, data gloves, or even body suits to some people. This thesis considers a virtual world to exist anytime that a three-dimensional environment exists where a user can navigate and interact with the 3D space. While some applications use the above devices to enhance the *immersion* of the user into the environment, for other applications these devices are unnecessary. NPSNET is an example of a relatively low-cost virtual world system which offers effective 3D navigation and interaction, using a spaceball as the only 3D-input device.

Virtual worlds can be defined as any hazardous, fictional, or simulated environment modeled by a computer system [Zyda91]. Although some authors treat the terms *virtual*

reality, *artificial reality*, *virtual world*, and *cyberspace* as having approximately the same meaning [Krueger91, Furness91, Rheingold91], the term *virtual world* is used herein to convey an environment that may be either real or unreal, embracing both *virtual reality* environments (which specifically attempt to represent *reality*, an actual environment) and *cyberspace* environments (which specifically attempt to represent unreal, imagined, or micro- or macro-environments¹).

C. EXAMPLES OF VIRTUAL WORLDS WITH HYPERTEXT

Many examples of both virtual world and hypertext applications exist separately. The idea of hypertext was probably first conceived by Vannevar Bush, as publicized in his historic article “As We May Think” in the *Atlantic Monthly* in 1945 [Bush45]. In this article, Bush proposed the Memex (memory expander) system, which contained many of the elements we today attribute to hypertext. Although the Memex system was never fully implemented, Bush’s writings provided the blueprint for future hypertext systems, and inspired future hypertext pioneer Doug Engelbart, who later invented the mouse and pioneered advances in interactive computing [Nielsen90, page 31]. Current advances in hypertext and hypermedia systems are being made on many fronts, including work that is funded by many different government and commercial sponsors at the MIT Media Lab.

Virtual worlds, as contrasted with hypertext systems, are an even more recent phenomenon. Workstations with real-time interactive 3D capability did not become widely available until the late 1980’s; indeed, two-dimensional graphical user interfaces only became widespread with the advent of Apple’s Macintosh in 1984. Although scientists and engineers were experimenting with graphical interfaces in the early 1970’s, primarily for simulation and training applications. The relatively high computational requirements of rendering three-dimensional objects and updating them in real time has, for the most part,

1. Micro-environments (environments containing enlarged representations of things smaller than can normally be seen with the unaided eye) and macro-environments (environments containing reduced representations of things that are normally very large) contain elements of both *cyberspace* and *virtual reality*, since they attempt to depict items that are real, but present representations that are normally unavailable to us.

limited virtual world modeling to either the IRIS workstations offered by Silicon Graphics or dedicated simulation systems, such as those developed by Evans and Sutherland [Rheingold91, p. 204].

1. The Information Visualizer

The Xerox Information Visualizer, developed at the Palo Alto Research Center (Xerox PARC), is one of the only known examples of a virtual world that incorporates a hypertext-like structure. The Information Visualizer reorganizes information into graphical objects that can be manipulated in 3D, using a metaphor called 3D/Rooms. The user is presented with a three-dimensional representation of a room, which contains familiar objects such as a table, a bookshelf, a chart on the wall, and one or more doors, along with unreal objects, such as buttons or icons. With the 3D interface, the user can navigate through the room, examining the various objects, and can select icons which present information related to their context. These icons include buttons attached to the floor, which flip up toward the user and present a variety of options when selected; textual items on a time line, which expands, contracts, and scrolls through time as selected by the user, and more objects which offer other interesting applications [Card91]. The user can then leave the room by navigating through one of the doors, and enter another room which has items grouped in a different context, perhaps for a different purpose altogether.

This work at PARC created a paradigm for what they called an Information Workspace, an extension of the desktop metaphor. An information workspace increased the amount of information directly available to the user by using 3D interface techniques to enhance the query tools available to the user. This system addressed the problem of database retrieval, but went beyond the simple retrieval of information to incorporate the actual use of the information once it is retrieved. Still, the virtual world exists in this system primarily for the purpose of more efficient and effective retrieval of the information encoded in the hypertext, by providing a mechanism to “visualize” the information. [Card91, Mackinlay91, Robertson91].

One of the data sets used by Card's group at PARC was a unix-like hierarchy tree that used three-dimensional modeling to wrap the lower tree levels around imaginary circles. This provides a carousel-like effect: when a file near the "back" is selected, the carousel rotates the hierarchy toward the user at and below the selected level. The PARC group found that if the rotation occurred instantly, without showing the intervening positions of the hierarchy while it was rotating, that the user became somewhat disoriented. The animation of the carousel, however, completely removed the disorientation. This observation led to a design decision for Hyper-NPSNET's navigation through portals, discussed below [Card91].

D. ADVANTAGES OF COMBINING VIRTUAL WORLDS AND HYPERTEXT

The ability to traverse a 3D space in a virtual world provides the user with the capability of interacting with virtual objects and exploring environments that depict either an actual or an imagined environment. Hypertext allows the user to navigate an information space in a non-sequential sequence, based on the user's own needs. Although virtual worlds by definition offer the ability to interact with the environment in real-time with no pre-defined navigational script, they generally restrict the user to traveling in real time.

Many virtual worlds offer only the ability to interact with the objects in the virtual world. It is not uncommon to see spheres or cubes in a virtual world which can be pushed or caught, responding with physically-based modeling as if they existed in a world with a prescribed amount of gravity, and the desired resilience [Rheingold91, Mandala91]. Some virtual worlds offer the user "walking tours" that allow the user to walk or bicycle on a special-purpose platform (while grasping a hand rail) to navigate around [Brooks91], and still others allow the use of virtual tools on virtual objects [Bolas91]. One virtual world even offers the ability to eat virtual food, ordered from a virtual waiter [Naimark91].

Partly due to the current limitations which exist in 3D interface design, the navigation of the virtual space can be quite tedious. For example, a virtual world that depicts a factory containing many different work areas may offer a useful simulation of manufacturing tasks,

but a supervisor in training may need to observe or explore several stations, perhaps at the same instant in time. Navigating in real time around the virtual space to locations that are well-known in advance is at best inefficient, and at worst unusable.

Many applications require more than simple interaction with virtual objects, particularly education and training applications. The virtual world becomes more of a setting for the actual work involved: familiar objects are represented as 3D icons that perform a function related to the context of their appearance. In many cases, the ability to examine information related to the objects, possibly while artificially frozen in time, provides for dramatically increased functionality. In NPSNET, the user needs to be able to select a vehicle and learn about its capabilities. After reviewing the information on the vehicle, the user could need amplifying information or information on related topics; for example, the user could want a list of vehicles that met specified parameters, or perhaps a list of vehicles both older and more modern than the chosen vehicle. Selecting an entire battle force, the user could need to review the historical locations of the force to determine the force's strategy and intended movements.

With an embedded hypertext, the user can not only request information on selected objects, but can also search for other objects that relate in a specified manner to the current object. If desired, the user can use the embedded hypertext links to recall this information. Perhaps more importantly, however, since each node in our 3D hypertext has a location in three-space, the user can "transport" directly to the location of a desired node's location in the virtual world. Using this capability, predefined locations, such as the ones desired in the factory example above, can be visited with a single command (e.g., a mouse click), eliminating the tedium of real-time interactive navigation through the virtual world. By taking advantage of these capabilities and merging the concepts of three-dimensional *cyberspace* and two-dimensional *hyperspace*, it is possible to create an environment that is many times more flexible and functional than either virtual worlds or hypertext can offer alone.

III. 3D/4D HYPERTEXT¹

As described previously, hypertext systems provide a means of cross-referencing related pieces of information. By devising links between related informational elements, the user of a hypertext can quickly access information that explains, enhances, or modifies the information currently at hand. Although the very first hypertext systems such as the Memex system visualized by Bush [Bush45] were intended to be paper-based, the current definition of hypertext implies the use of a computer [Nielsen90].

Shneiderman presents three “golden rules” that determine if an application is suitable for hypertext, and Nielsen adds a fourth [Nielsen90, page 43] [Shneiderman89]:

- A large body of information is organized into numerous fragments.
- The fragments relate to each other.
- The user needs only a small fraction at any time.
- Do not use hypertext if the application requires the user to be away from the computer.

These guidelines imply that some systems are not well served by hypertext implementation, which certainly is the case. In the case of NPSNET or most virtual world systems that have a need for lookup referencing, however, all of the above guidelines are easily met. The following sections describe some of the functionality that would be provided by an embedded hypertext in a virtual world application.

A. 3D/4D HYPERTEXT DEFINED

A suggested paradigm for a 3D Hypertext follows:

A 3D hypertext can be defined as a hypertext which meets the following two properties:

- Each hypertext node, or anchor, has a distinct location in three-dimensional

1. Throughout this thesis, the terms “hypertext” and “hypermedia” are used interchangeably. It is assumed that a hypertext will not be limited by design to nodes that are comprised only of text, but can also contain pictures, graphics, sound, animation, video, or some or all of the above.

space.

- Any hypertext node can directly access another node and its 3D location.

One added characteristic to the above results in a 4D hypertext:

- Hypertext nodes can be automatically updated over time, or can contain links to other nodes that have a specific temporal reference.

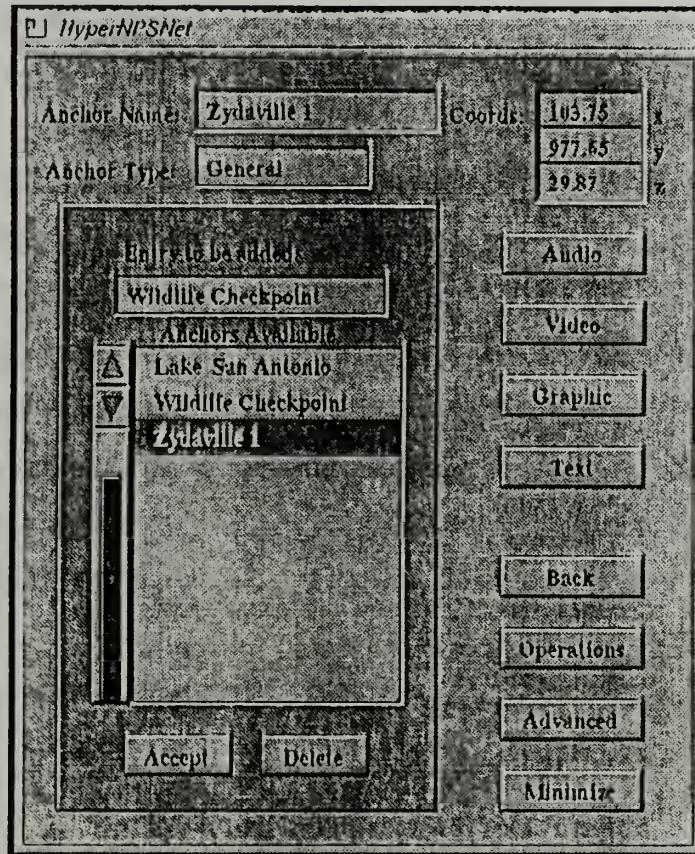


Figure 1.2 : Prototype of Hyper-NPSNET Interface

B. LOCATION OF HYPERTEXT NODES IN 3-SPACE

The fact that each node has a distinct location in 3-space provides the required functionality of lookup reference for any item, or any location, in the virtual world. That is,

at any given location in the virtual world, a hypertext anchor can be embedded in the form of a *3D icon*: an icon that has visible width to the user in all three dimensions, and has a unique identification number and noun name. Once this icon is selected, the available hypermedia elements are presented to the user, as illustrated in Figure 2.

These elements can be comprised of a simple text field, but will generally contain other hypertext nodes that provide additional options. For example, selecting a book icon entitled “History” could lead to options that include temporally-oriented text, animated sequences, a recorded quote by a noted historian, or amplifying graphics.

Additional functionality can be gained by locating the hypertext information at designated locations in three-space. At first glance, it may seem unimportant where in the virtual world the information is stored; however, by grouping the information at a specified location, both the user and the system can make both logical and efficient transitions from one area to another². One approach would be to construct a library, visible in the virtual world, that would contain all the supporting text nodes, sorted by subject, time, or other desired indexing method. Additionally, the textual node can, once referenced, provide links to other hypertext nodes (of any hypermedia type), which themselves have a relevant 3-space location.

C. SELECTION OF 3D HYPERTEXT NODES

As described above, in a 3D hypertext, each node is made visible by a 3D icon. Other non-unique attributes, such as media type, a noun name, a time flag for creation and modification, and other descriptive information can also be attached to each node to add flexibility to selection processes. The 3D nature of the identifying icon is vital to prevent it from being “invisible” or hidden to the user. The visibility of all icons can be toggled on or off at the discretion of the user, to prevent clutter effects in the virtual world, but once 3D

2. The ancient Greeks used spatial locations—buildings, walls, and so forth—as an aid to memory when recalling speeches. Hyper-NPSNET can, if desired, embed the *information* into the desired location, using either icons or the actual text.

hypertext is made active, any icon that is visible to the user can be accessed with the supplied pointing device, generally the mouse.

Icons are also temporally referenced; that is, they appear only when *in time* they are applicable. An application that makes good use of this characteristic is a wartime battlefield. Enemy forces arrive in, depart from, and cease to exist in the virtual world (geographically a subset of the enemy country, for example) at various times, so the hypertext icons that relate to them can be linked directly to the virtual objects themselves. Once a tank explodes, for example, its supporting hypertext nodes cease to exist, although other nodes not currently visible may continue to exist in an allocated 3D location to provide historical analysis, replay, and other applications. Due to the temporal nature of Hyper-NPSNET icons, color- or shape-coding is available to distinguish historical, real-time, and projected objects. Use of these icons can also prevent object clutter in the virtual world: the tank's actual position is represented by the tank object; turning on historical hypertext displays icons along the tank's actual track; and turning on projected hypertext displays the projected path of the tank according to intelligence estimates.

An important design consideration is selection ambiguity for the 3D icons. Since an icon exists in three dimensions but is selected by a two-dimensional device, it is possible to accidentally select an icon that is collinear in 3-space to the desired icon. To resolve this ambiguity, the hypertext node nearest the location that corresponds with the user blinks, indicating its proximity. Any other hypertext node, visible or not, can be accessible by alternative selection methods, described below.

In any virtual world application that uses one or more two-dimensional CRTs, the use of the various areas of the screen is an important interface decision. While most of the screen "real estate" will be used to display the virtual world, other areas (such as a specified vertical strip adjacent to the virtual world display) are generally needed to provide navigation control, help, or amplifying data options to the user that is navigating the three-space. In NPSNET, a horizontal strip at the top of the display gives information about the vehicle being operated by the user at the console, including fuel and weapons remaining,

velocity, and location. Display options are also available at all times to the user, such as an overhead “geographic plot” of the area which shows where the user’s vehicle is operating, and the location and description of other vehicles in the same area.

Some area of the CRT is required to display hypertext options. Optimally, the hypertext control panel should be displayed in a manner which would not obscure the view of the virtual world, so as to not limit the functionality of the user in the virtual world while using the hypertext, unless the chosen hypertext option is specifically designed to manipulate the virtual world. Macintosh systems, for example, directly support the use of multiple monitors, which can be used to display the desired information. In HyperNPSNET, however, multiple windows are used on a single screen to present a type of information consistent within each window.

One illustration of the benefits of a dedicated hypertext display area is direct selection of a desired hypertext node. While graphical selections of hypertext icons will be limited to those icons which are visible, the user could directly select a desired hypertext node by specifying its unique noun name or identification number. Further, a database-like front-end to the hypertext would provide the user with the ability to query the system for a list of available hypertext nodes that meet certain criteria specified by the user, such as all nodes in a specified proximity (see 3D zones, below), all nodes of a certain type, or all nodes that relate to a specified subject. Once the results of the query are presented to the user, the user can select any of the displayed nodes.

D. HYPERSEARCH

With the above numbering and location scheme, options can exist to provide the capability to search from the current selected node for similar or related nodes in the virtual world, providing the equivalent of a database front-end to access related elements, even those of diverse media. This type of search, which will be referred to as *HyperSearch*, could be based on diverse search queries, including search for animations on the selected subject, text files that relate to the current node’s descriptive attributes, or general searches for all

node types within a specified time frame. HyperSearch gives the user, and the virtual world planner, options which significantly extend the functionality of applications in training, analysis, modeling, and other applications.

1. 3D Zones

A supporting metaphor for HyperSearch is a data structure called *3D zones*. These zones exist invisibly to the virtual world in normal use, and serve as a location grid that helps provide spatial orientation to hypertext objects in the virtual world. A lookup table provides cross-referencing between virtual objects and their corresponding 3D zones, preventing the user from the unnecessary chore of annotating which zones correspond with objects of interest. Each 3D zone can have attached fields which identify the principle object or group of objects that exists in the zone. Searches for hypertext objects can be constrained to areas within a specified set of 3D zones.

a. Types of 3D Zones

There are two types of 3D zones: fixed and moving. Fixed 3D zones are analogous to latitude/longitude/height-or-depth grids, and are used as positional references in the virtual world. Moving 3D zones are attached to principal objects in the virtual world. These moving 3D zones have a parameter which corresponds to a three-dimensional offset from the object to which they are attached. As an example, a moving 3D zone can be created to correspond to a 2000-kilometer offset from the center of an asteroid. A HyperSearch can be performed to determine objects that exist, at a specified time, within the 3D zone. A moving 3D zone exists at all times within exactly one fixed 3D zone, which specifies the location of the moving 3D zone in the virtual world.

b. Type of Applications Utilizing 3D Zones

In addition to obvious correlation with virtual reality applications (latitude, longitude, and height or depth), 3D zones have interesting possibilities in *micro-worlds* (virtual worlds where normally small objects have been greatly enlarged, such as a

visualization of the inside of a molecule) and in *macro-worlds* (such as a view of the solar system as viewed from the galaxy Andromeda). In a macro-world of the solar system, for example, the user can perform HyperSearch for all nodes of type picture that exist relating to the 3D zone corresponding to the planet Neptune. Additionally, some applications of virtual reality make good use of 3D zones. Such an application is detailed in section (d) below.

c. Flexible Hypertext Links

Since objects in virtual worlds (in most applications) move, hypertext links to those objects must also move. This has significant implications to the architecture of the hypertext, which can be envisioned as a web-like structure with joints and feet that are constantly moving. To maintain an accurate structure of this type is certainly expensive from a computational point of view.

However, due to the object-oriented nature of Hyper-NPSNET's architecture, the hypertext links may be formed dynamically; that is, when they are requested by the user. The moving object's position is updated over time as it is currently with NPSNET. When the user enables hypertext visualization, the position of hypertext links associated with objects selected by the user are updated.

d. An Application of 3D Zones

Some applications require flexibility when creating moving 3D zones for virtual objects. Certain objects can be best associated with zones by shrinking the 3D zones to fit the objects, which results in 3D zones with various shapes (spherical for planets, rectangular solids for navigational airspace, irregular for asteroids and bacteria, for example). The actual shape of the 3D zone may or may not correspond directly with the shape of the object.

An example that illustrates the importance of 3D zone shape irregularity is a submarine's moving haven. In open ocean, a submarine is generally required to stay within a moving haven, or MHN, which corresponds to a rectangular solid that moves at a

specified course and speed. This arrangement serves to deconflict interference with other submarines and surface shipping.

For our example, consider the MHN as being 50 NM (nautical miles) by 10 NM, from the surface to a depth of 150 feet, traveling on a course of 045 degrees true at 12 knots. The submarine must remain inside this MHN at all times until the MHN changes, in accordance with the PATORD (patrol order) issued by the submarine's operational authority. To calculate a moving 3D zone that precisely delimits the MHN to every nuance of the submarine's curvature is to subject the system to unnecessary and impractical precision, in this case. An MHN that forms around the geometric center of the submarine is more than adequate. However, when the submarine is navigating restricted waters, such as when egressing from a port with a twisting, shoaling channel, it is imperative that the submarines "envelope of safety"—the distance from any point on the hull to the nearest shoaling—be accurate to a matter of feet. A 3D zone constructed around a virtual submarine to train the navigational team needs to be shaped with a high degree of accuracy.

2. An Application of HyperSearch

Consider an intelligent tutoring system application of Hyper-NPSNET that would provide training for aerial combat maneuvers. The user, flying the simulator in real time, can identify a distant airborne target, freeze the simulation, and use HyperSearch to bring up information on the target type design, numbers that currently exist in the enemy inventory, current known deployment patterns (temporally-significant data), and other desired data. The user could then enter a mock-up of the enemy aircraft (which would be a virtual object within the virtual world, or a *meta-virtual object*) and analyze the aircraft's capabilities from the perspective of the enemy cockpit. An animation of intelligence films of actual enemy aircraft maneuvers could be retrieved and examined using HyperSearch, as well as surveillance photographs of actual enemy aircraft. After reviewing the desired information on the target, the user can continue the simulation from the point at which it stopped.

The ability to not only review information, of varied media types, that relates to a specified object, but to also process information that is gathered *temporally*, *spatially*, and *dynamically* about an object is a valuable analysis tool. A continuous search can be performed in the background, informing the user when new objects of a specified type enter the specified object's 3D zone envelope. A proximity search on a hemoglobin cell can reveal the presence or absence of cancer cells in a micro-reality system. HyperSearch offers the user tremendous power in the use of a virtual world.

E. HYPERNAVIGATION

1. Hypernavigation Defined

Another exciting capability of Hyper-NPSNET's implementation of 3D/4D hypertext is the ability to physically traverse the 3D hypertext links to the location of a selected hypertext node in the virtual world. Hypernavigation can be defined as the movement of the user through the virtual world to a location specified by the selected hypertext node. Through such navigation, the user can visit specified places in the virtual world by direct access, instead of being required to navigate the virtual world in real time to a location that is already well-known.

2. Portals

Hypernavigation greatly extends the power of the user in the virtual world, and enhances the immersion of the user into the virtual environment. Any hypertext node is now a portal, through which the user can pass on his or her way to a specified location. The capability described above to directly access a hypertext node, by specifying its name or its identification number, is essential to the implementation of hypernavigation. As an example, a user could specify an object's 3D zone, and immediately be transported to the nearest boundary. With the database front-end capability provided by HyperSearch, the user can view the results of a query, such as "display all virtual world objects meeting the following criteria:...", *and then visit each one of them.*

3. The “Alice” Effect

Work done on the Information Visualizer at Xerox PARC led to the observation that the user can become disoriented if an object is instantly transported to the user. Even in an example as visually consistent as the view of a unix hierarchy, the instant rotation of a carousel-like sublevel added confusion to the user. This disorientation was eliminated by providing animation of the carousel at a rate that was easily perceptible by the user, but still did not cause undue selection delay of the desired object [Card91].

Hypernavigation literally adds many dimensions to this problem. In a virtual world, the user can be instantly transported to a new location in 3-space, sometimes known as the “Star Trek transporter” or “Alice” effect.³ Elimination of this disorientation is reduced, if not eliminated, by enabling the user to literally ride the hypertext links to the location of the selected node in the hypertext. The speed at which the user travels along the hypertext link is proportional to the distance that the user must travel to arrive there, and can be adjusted by the user. In the factory model, for example, the user may want to go from the prefitting station to the frame assembly station rather slowly, so as to not become disoriented while navigating the factory space. This is facilitated in the design of the factory by presetting a rate of travel, applicable only while inside the factory, to an appropriate rate. However, a trip to the virtual library, which exists on the roof of the factory in 3-space, may occur at near instantaneous speed, since the virtual library only exists in cyberspace whereas the factory exists in virtual reality. After visiting the virtual library to review reference materials (or perhaps the performance history of employees who work at that station), the user can rapidly return to the exact departing location in the virtual reality of the factory. After returning to the factory, the *Hypernavigation rate* returns to the slower value preset for the factory.

3. “(Hypernavigation can lead to an effect)...not unlike Alice (in Wonderland) falling down the rabbit hole, by the way.” - Dr. Michael Zyda, March 1992.

4. Hypernavigation Interface

Advances in interface capability can extend the power of hypernavigation. Ideally, a voice-recognition system would allow the user to interface directly with the hypertext control panel, providing the ability for the user to traverse the virtual world quickly and efficiently. A voice-recognition interface for HyperSearch, for example, can enable the user to issue voice commands for the destination desired, or specify parameters for HyperSearch. Once HyperSearch finds a valid result to the search parameters, the user can be immediately transported along the 3-dimensional hypertext links to the destination of a specified node in the search product.

As effective as a voice-recognition interface can be, there are some environments that preclude the use of such an interface (sp., those with extremely high levels and diverse frequencies of background noise). Hypernavigation is in no way predicated on high-level interface implementations, however. The user can specify the location desired using an interface as primitive as a command-line prompt system⁴, for example, and be transported there as soon as the enter or return key is struck.

5. Three-dimensional traversal tracking

In addition to traversing a virtual world via hypernavigation, the hypertext can also be designed to enable the user to backtrack along the path of traversal. A simple voice command, button selection, or typed command takes the user to the previously-visited node, or allows travel in a reverse direction on the hypertext link until stopped by user input. This would allow the user to literally backtrack in three dimensions, which would provide important functionality for simulation scenarios where temporal analysis is desired.

As an example, in a virtual world modeling air-to-air combat, it is useful to reconstruct the scenario of close-in air-to-air combat (dogfighting) and analyze the position

4. This example should in no way be misinterpreted as a recommendation on the part of the author that such an interface be used, especially in conjunction with system such as Hyper-NPSNET.

of the user's aircraft as a specific point in time relative to opposing forces. The hypertext can be designed to "self-author" by "dropping" hypertext nodes periodically along the flight path of the user's and opponent's aircraft and anchor themselves in 3-space, providing a convenient mechanism for replay.

F. INSTANT ASPECT CHANGE

Since hypernavigation makes possible direct traversal to any specified location in 3-space, the virtual world can be designed to default to a grid of hypertext nodes which exist around any desired object or designated sets of objects. In a tactical battle simulation, for example, a diamond-shaped grid of nodes can be constructed around each air or ground platform in the virtual world. By directly specifying the location, such as "above", "behind", "below", "front right", or "rear left", for example, the user can jump out of the current vehicle and view it from the desired perspective. This would provide the ability of the user to directly view the current location in the virtual world from a variety of perspectives, including the frequently-implemented "God's eye" view. By specifying another vehicle in addition to the location on the grid, the user can be transported to varied perspectives of other vehicles in the virtual world, including an opponent's perspective, to support such objectives as aspect analysis for scenarios such as air-to-air or submarine combat scenarios.

An example application sensitive to aspect change is a proposed extension of Hyper-NPSNET to provide submarine tracking simulations. Since the sound frequencies radiating from a submarine are extremely aspect dependent, and the target's course must be accurately approximated for weapons employment, it is extremely important to determine the aspect of the target. A Hyper-NPSNET-based intelligent tutoring system could display frequencies detected by sensors from the target, and then provide the user with the capability to view the enemy submarine's actual aspect. Additionally, the user could jump into the enemy submarine and view himself and his radiated frequencies, and perform vulnerability analysis. After the training scenario is complete, the user can return to the

location and time of critical portions of the engagement, and complete the scenario with alternate actions to analyze the effectiveness of tactical options.

IV. AUTHORING IN HYPER-NPSNET

A. THE NEED FOR AGILE AUTHORING

Key to the success of Hyper-NPSNET is the ability to efficiently and effectively modify existing virtual worlds or create new ones. Many virtual worlds require updating to ensure accuracy and/or realism, since they contain temporally-sensitive objects or environments whose rapid change requires the addition and deletion of objects. Some virtual worlds require modification by the user to provide for more effective interface paradigms; for example, intelligent tutoring systems that must adapt to the needs of the user or group being trained. Hyper-NPSNET will have advanced facilities for adapting to changing environments.

Hyper-NPSNET's authoring system will be usable on two levels:

- User level: high-level updating and modification functions
- Design level: high- and low-level authoring features

B. OBJECT-ORIENTED CONSTRUCTION

Hyper-NPSNET's object-oriented construction facilitates many of the high-level authoring features planned for use by many applications. Objects in the virtual world will be *objects* in the object-oriented programming (OOP) sense, and the construction of new objects will be facilitated by the ability to inherit the characteristics of existing objects, providing for maximum reusability of code. Behaviors of the objects that exist in the current virtual world can be inherited selectively when creating new objects to provide for specialized types; for example, "Apache", a new instance of class "helicopter", can be created that has different speed, handling, weapons, and sensor characteristics, but in all other respects resembles the generic class "helicopter".

C. USER-LEVEL AUTHORIZING

Hyper-NPSNET's user-level authoring provides the capability to perform rapid and effective modifications to the virtual world. It is important that the user be able to efficiently modify the battlefield in wartime, for example, to account for the destruction of objects, relocation of objects already in the environment, and the addition of newly-deploying forces. A static and rigid virtual world would be completely useless as a real-time analysis tool, and is only suitable in predictable training situations. Since wartime is neither static nor predictable, Hyper-NPSNET's agility serves a major part in its functionality.

In the user-level authoring mode, the user is able to select an object and create a new instance of the selected object with a simple command from the keyboard or another input device. Existing objects can be repositioned by a simple "move" function, which provides the ability to input the exact new location of the object, or the option to select the object and accelerate it in a specified direction and speed.

The user can also assign hypertext links from and to objects. As all hypertext links are bidirectional, the user can simply specify two objects and assign a link which connects them. Location and specification of the desired objects can be accomplished using many of the tools described above, including HyperSearch and 3D Zones. If desired, the user can perform a HyperSearch and connect all nodes in the search product to each other, making a K_n -type graph structure.

If necessary or desired, the user can add or modify objects that exist in the virtual library. Some applications require nearly constant updating of data files, especially those in highly dynamic environments. Hyper-NPSNET can give the user access to any unprotected nodes in the virtual library, including text files, pictures, animations, and sounds.

Protection of nodes is performed at the Design Level. Any node with a protection flag set can only be modified by a user with a Design-Level logon verification.

In some but not all applications, the user can add multimedia node objects, such as pictures, sound, and animation. Although most applications do not require the user to manipulate objects of such diversity, others, such as Intelligent Tutoring System design, multimedia-based research (such as analysis of language dialects or natural language processing), and other applications require the user to become deeply involved in one or more multimedia elements. Hyper-NPSNET provides the capability for the user to manipulate these objects without requiring the user to learn a text-based computer language. (As more workstations become available with digital signal processing (DSP) and built-in CD-ROM capability, the provisions in the system to manipulate graphics, animations, and sound increase rapidly.)

D. DESIGN-LEVEL AUTHORIZING

In addition to the functions provided in User Authoring, the virtual world designer has many other tools with which to design and modify virtual worlds. The assumption implicit in the design of the User Level is that the virtual world will not be modifiable to the user, but the objects in the virtual world will. Therefore, the Designer can essentially move the walls of the virtual world in any direction; add, delete, or modify 3D Zones, adjust preset values for Hyper-navigation speeds, and perform reliability and interface checks (including but not limited to a check for nodes that have no links assigned).

Designers also have the capability to assign and modify parameters assigned to structures in the hypertext, such as the labels or offsets for 3D Zones. Most of these functions for design-level authoring are supported by high-level tools and structures, which precludes the need for entering of massive cryptic commands to accomplish a small modification to the environment. Additionally, users can modify the interface panel of applications that were created by them, allowing the opportunity to tailor the interface to the actual users of the system.

In some cases, the virtual world requires connection to remote databases to update the status of objects in the hypertext. These functions are performed by those with Design Level authoring.

E. AN EXAMPLE OF AGILE AUTHORIZING

In modern U.S. Naval submarine simulators, a battle stations team can detect targets using the same sonar gear as is installed on the submarine, employ simulated weapons, and perform tactical maneuvers similar to those performed in battle. The target is generally a simulated ship or submarine that radiates noise levels by accessing digitized files of actual targets, then adjusting the noise levels to take into account doppler and other effects [Daley89].

One of the difficulties in operating a trainer such as the one described above is validation of the simulated targets. Since enemy naval capabilities change rapidly, there is nearly a constant need to modify the target simulations to provide for new weapons capability, changes in the radiated noise from each target, and structural modifications to ship classes that change maneuvering and acceleration parameters. Typically, the process to add one new submerged target to the system can take years, and thousands of dollars, to accomplish.

In a system such as Hyper-NPSNET, the user has the ability to add or modify targets with the above high-level tools. While modification of sound parameters is not expected to be trivial, the user can modify the handling characteristics, and possibly some of the weapons parameters, by using the high-level scripting tools provided to manipulate the object files. Once these modifications are complete, the target can be validated by an intelligence gathering team or other suitable personnel.

V. APPLICATIONS OF HYPER-NPSNET

A. CHARACTERISTICS OF SUITABLE APPLICATIONS

As illustrated in the previous pages, there are many applications for virtual worlds with embedded 3D/4D hypertexts. It is illustrative, however, to provide an extended scenario in which Hyper-NPSNET provides exceptional functionality.

Before embarking on an extended scenario, however, the following assumptions are submitted as entering arguments in considering the use of Hyper-NPSNET:

- The application's domain is in an environment which has three-dimensional or four-dimensional (e.g., 3D plus temporal) qualities.
- Graphical workstations with sufficient rendering power, such as the Silicon Graphics IRIS series used with Hyper-NPSNET, are either available or affordable for the application.
- The user is capable of manipulating the 3D interface.¹
- Sufficient assets are available to update the virtual world as temporally-sensitive data changes, or other changes occur in the virtual world.

1. Fighting the Virtual Battle

Development of tactics for a planned or possible tactical battle scenario is critical to effective battle planning. Such simulations can reveal necessary changes in required armament, force strengths, deployment and resupply strategies, and other factors.

To be effective for military battle planning, a simulation system must present realistic depictions of the environment, provide agile authoring which is robust enough to provide for fast and efficient updating of the environment, offer effective real-time interaction with both the environment and the simulated opponent's forces, and provide powerful analysis tools to the user. While many current simulation systems offer realistic depiction of the environment, and perhaps the opposing forces, few offer the ability to

1. To accommodate users of various computer experience, as well as handicapped users, interface design is critical. Alternate input devices, including the use of voice input, can provide for the accommodation of users of various levels of computer sophistication and physical attributes.

engage the opponent in real time, and none offer the capability of performing this interaction on a low-cost workstation. Moreover, no other known simulation system provides the agility necessary to modify the environment or elements within it to adapt to rapidly-changing conditions found in a typical battle scenario.

Battle-planning simulation begins with a depiction of the geographical environment; for example, an area near the Iraq-Kuwait border. Since the environment being simulated has significant impact on the performance of the system, size and resolution decisions are made that are tailored on the system's intended use. The size of the area being modeled is dependent on the scope of the user: a commanding general requires a significantly larger area than does a battalion commander, for example, and is interested primarily in terrain and enemy force simulation. An aviation simulation, on the other hand, has broader space and different resolution requirements than does a ground-attack simulation.

Once the environment has been modeled to an acceptable level of detail, friendly and enemy forces are placed in the desired positions. In a hypothetical scenario, friendly forces are positioned in the area of most probable staging, and enemy forces are placed either in currently-known locations or in areas where they would likely be found in the scenario being simulated. In actual wartime, the ability to update the actual position of both friendly and enemy forces is crucial to maintaining a reasonably accurate depiction of the virtual reality being modeled. In support of this requirement, the Design Level author provides an automatic updating capability that reads in both targets and friendlies by identification number, then updates their positions in the virtual world. The User Level author now only has to verify the correctness of the data which is automatically input. (Target locations can also be updated manually, if desired.)

The current NPSNET prototype [Zyda91] offers the capability of autonomous and semi-autonomous forces. Hyper-NPSNET builds on this capability by providing temporally-referenced analysis, which project the current virtual reality into a variety of scenarios in cyberspace. The various cyberspace environments, which correspond to

alternative tactics by both the enemy and friendly forces, can be visited, analyzed, catalogued, and evaluated, to determine the outcome of alternate strategies. Enemy positions can be visited to analyze the battlefield perspective from the enemy point-of-view. Enemy weapons can be simulated against friendly forces during each scenario and analyzed for effectiveness. Air missions can be simulated against the enemy encampments to analyze for possible vulnerability. The current virtual reality can be analyzed using HyperSearch to determine the strength of enemy forces, and 3D Zones can be established to provide warning points when specified proximity boundaries are breached.

The ability to model such a scenario on a system such as Hyper-NPSNET adds realism, agile authoring, and powerful analysis capability compared to systems currently used for these purposes.

2. Magnetic Disk File Maintenance Training

Learning to perform maintenance on a device such as a magnetic disk file (MDF) is frequently complicated by the fact that very few exist outside of shipborne platforms. On FBM submarines where two crews operate the submarine, for example, the crew ashore must undergo constant training to ensure readiness while the opposite crew is on patrol. The availability of many shipboard systems for training is poor to non-existent.

Effective training on such systems is facilitated by three-dimensional modeling in the Hyper-NPSNET virtual environment. In the virtual world, the user can interact with the components parts of the hardware using button and knob controls, as well as pointing devices such as a mouse or spaceball. Combining the virtual modeling with a hypertext-based Intelligent Tutoring System provides the user with the capability of reference lookup, what-if analysis, hypernavigation, and hyper-search.

Maintenance routines can be simulated by providing the trainee with a textual representation of the required procedure, along with a virtual toolkit. This toolkit contains the required manual tools, electronic devices such as a voltmeter, and safety devices such as electrically-insulated gloves. The trainee positions the probes of the multimeter on the

test terminals, and the voltage or amperage that results is displayed on the virtual multimeter. If the trainee touches a live circuit, appropriate simulation of an electrical shock results, followed by training on electrical safety retrieved automatically from the virtual library's technical reference. If a fire breaks out, the user is presented with appropriate virtual damage control gear with which to combat the casualty.

During one of the maintenance procedures, the trainee becomes confused as to the operation of the electro-hydraulic system which drives the MDF read/write heads. By travelling down the hypertext link to the hydraulic system, the user enters a micro-world where he is surrounded by the electro-hydraulic system. If desired, the user can become a passenger in the hydraulic system, following the path of the fluid as it is forced from place to place, then repeating the journey with specified isolation valves shut.

In some cases, there is a thin line between Intelligent Tutoring Systems and Expert Systems. The same knowledge base that is used for training can also be used to guide users that perform maintenance procedures, for example. The structure of the underlying hypertext provides the application designer with a convenient vehicle for expert system development.

3. DSRV Piloting

Learning how to operate a Deep Submergence Rescue Vehicle (DSRV) to the precision required for docking with a sunken submarine, in amongst strong ocean currents, is a formidable training task. By using many of Hyper-NPSNET's features, including instant aspect change, hypernavigation, and hyper-search, the user can begin to learn to navigate the DSRV without actually using the vehicle, saving costly time and man-hours.

In the authoring phase, the DSRV and the sunken submarine are modeled in the virtual world. Handling characteristics of the DSRV are implemented and certified to be correct, and the actual mechanical devices used to control the DSRV are simulated. Ocean current modeling is performed, based on actual data from the locale of the simulated

emergency. The DSRV takes advantage of physically-based modeling to provide realistic buoyancy and turning characteristics.

Hyper-NPSNET's Intelligent Tutoring System to train users in DSRV piloting techniques can be set up to record progress of each specific user, including areas of difficulty or mastery. Additionally, casualties can be programmed into the simulation, to train on the proper response actions for fire, flooding, collision, etc.

VI. SUMMARY AND CONCLUSIONS

This thesis presents a blueprint of a system yet to be, a marriage between a dynamic, user-adaptable way of organizing information, and the most revolutionary step yet achieved in interface design. There are many issues that remain that will shape the implementation of systems such as Hyper-NPSNET. Some of them are presented below.

A. VIRTUAL WORLD DEVELOPMENT

Due to the relatively recent emergence of virtual world capability on low-cost workstations, it is somewhat problematic to make assumptions about what direction virtual world development is taking. Some still think it's too expensive; others are disillusioned by scratchy, jerky demonstrations, although important advancements are being made at improving interactive display smoothness without increasing CPU power [Teller91]. While some see virtual worlds as an important new capability for advanced computing, still others live in fear that virtual worlds will become the escape of choice for future generations, surpassing television as a way to "immerse" oneself in a distracting world, or even providing the gateway to the development of cyborgs and other machine-enhanced creatures [Foley87, Gibson84, 1Gibson86, 2Gibson86, Gibson88].

Any powerful tool, be it a workstation or a buzz saw, offers the potential for abuse, and virtual worlds are no exception. Although one would hope that most users will use the tools with care and appreciation, there will always be those few that attempt to use it for illegal and/or selfish reasons.

The above is, however, no reason to deprive the rest of us from a powerful, affordable, and effective tool. Virtual worlds can present environments to the user with a realism matched by no other system. Objects can be rendered in their familiar three-dimensional context, just like we are used to in the actual world. Creatures as complex as articulated beings can already be manipulated in very life-like simulations [Badler91]. Once modeled,

many virtual objects can be manipulated in 3-space using simple and cost-efficient input devices. When a sufficient number of these devices have been made and used in a variety of applications, it will be up to the evangelists of HyperWorlds to increase the accuracy of the public's perception of virtual world systems, and to bring the message of the cause to the system designers of the world [Kawasaki91].

B. MODERN HYPERTEXT TRENDS

Hypertext is finding an increasing number of applications for an increasingly large base of computer users. Systems such as Apple Computer's HyperCard, given away free with every Macintosh sold, have provided a large base of end users and a surprising number of hypertext designers[Apple91, Goodman90]. Shareware archives have gigabytes of hypertext applications, written programmers ranging in expertise from novice to expert programmers. Nearly every variety of personal computer has a hypertext language that provides convenient high-level hypertext design. Some applications use hypertext as an ideal on-line help system.

C. VIRTUAL WORLD INTERFACE ISSUES

Some applications require the full immersion capability of 3D goggles, datagloves or bodysuit for effective interface with the virtual environment [Kelly89, Lanier91], Many more, however, require only display interfaces such as that used for Hyper-NPSNET (2D display, mouse, spaceball, knobs, etc.). As with many questions regarding the use of computer systems, there is no absolute answer as to the type of interface that should be applied for all applications.

As the complexity of the virtual environment and the underlying hypertext, increase, the need for more sophisticated interface technology is inevitable. Just as there was no apparent need for the mouse as an input device a short time ago, and a spaceball was not envisioned by most, tomorrow's Hyper-NPSNET-like systems will require the use of higher abstractions in the use of interface devices.

Voice-recognition devices, considered impractical and overly expensive just a few years ago, are now finding their way into the commercial marketplace. Apple Computer is already demonstrating a voice-recognition system that runs on a Macintosh (nearly all models of which come with a built-in microphone), NeXT computers have voice e-mail capability, and Silicon Graphics' IRIS Indigo includes built-in digital signal processing (DSP) [Rheingold85,Hollis92,Stahr91]. Commercially-available voice recognition systems exist now for personal computers, and have enjoyed some degree of success, as evidenced by the competition starting to present itself in the marketplace [Hollis92].

One company, Bright Star Technologies, markets a low-cost product that models a human agent that appears on the screen and performs requested functions [Gasper91]. Although most of the interface with the agent is by mouse selection, a limited voice interface exists, and can be further developed as personal computers develop increasing voice recognition capability.

In virtual environments, agents can assist with many of the functions found to be cumbersome in current virtual world systems. An agent that interprets either a voice or tactile input to increase or decrease the speed at which the user travels through the virtual world, for example, gives the user more freedom to explore the virtual space. Providing the user with the ability to name the 3D Zone which corresponds to the desired destination, for example, provides tremendous flexibility and efficiency. Virtual agents with natural-language interfaces may be the ultimate front-end to HyperSearch: few interfaces can improve on the efficiency of interactive dialog when searching for a particular bit of information [Brand88]. Whether the agent takes the form of a human or not is an interface issue; the resulting functionality is a major design issue in the architecture of increasingly-complex systems [Laurel90].

D. FOUR-DIMENSIONAL HYPERTEXT

The development of virtual worlds is still in its infancy. Currently, much exploration into virtual worlds is severely limited by the power of the average workstation processor,

and many people still think of virtual worlds as a drug-induced psychosis used to escape the real world. Clearly there is a lot of education in the future for virtual world advocacy.

With recent developments in the engineering and the marketing of 3D workstations, however, virtual world development is poised on the point of no return. With aggressive virtual world development also comes the reaction from the general public, which can range from indignant self-righteousness to paranoia to a complete embracing of the technology.

Hyper-NPSNET makes the point that hypertext is ready to go not 3-D, but 4-D. With virtual systems that are referenced to an instant in time—either past, present, or future—a hypertext can traverse time and grow or shrink appropriately. More than just simple retrieval of text, a hypertext in a virtual world can maintain a database that describes, even circumscribes, the environment being modeled in the virtual world. Powerful tools such as Hyper-Navigation, HyperSearch, and 3D Zones greatly enhance the functionality of the integrated system. With its powerful analysis capabilities and affordable cost, Hyper-NPSNET can write a new chapter in the organization and presentation of information, and in the development of HyperWorlds for diverse applications.

REFERENCES

- [Apple91] Apple Computer, Inc. “(System 7) *HyperCard Basics*”, 1990-1991.
- [Atwood91] Atwood, Donald J., *DoD Modeling and Simulation Management Plan* (Memorandum from the U.S. Deputy Secretary of Defense), ser. 38917, 21 June 1991.
- [Badler91] Badler, Norman I., Barsky, Brian, and Zeltzer, David, eds. “Task-level Graphical Simulation: Abstraction, Representation, and Control” by David Zeltzer. *Making Them Move*. (Chapter 1). San Mateo, California: Morgan Kaufmann Publishers, Inc., 1991, pp. 3 ff.
- [Bolas91] Bolas, Mark. “Be Here Now” (Fake Space Labs), a virtual world exhibit demonstrated at ACM SIGGRAPH ‘91. *Computer Graphics*, Vol. 25, Number 4, July 1991, p. 382.
- [Brand88] Brand, Stewart, *The Media Lab*. New York: Penguin Books, 1988.
- [Brooks91] Brooks, Fred, Fuchs, Henry, and Robinett, Warren. “Five Applications using Head-Mounted Displays” (University of North Carolina at Chapel Hill), virtual world exhibits demonstrated at ACM SIGGRAPH ‘91. *Computer Graphics*, Vol. 25, Number 4, July 1991, p. 384-5.
- [Bush45] Bush, Vannevar, “As We May Think”. *Atlantic Monthly*, July 1945, pp. 101-108.
- [Card91] Card, Stuart K., et. al. “The Information Visualizer: An Information Workspace”. *Human Factors In Computing Systems* (ACM SIGCHI Conference Proceedings), 1991, pp. 181-8.
- [Daley85] Personal experience by the author as cruise missile and fire control system instructor at Naval Guided Missile School Detachment Norfolk, 1985 - 1987, and as strategic weapons officer on two U.S. Naval submarines, 1987 - 1989.
- [Daley89] Personal experience by the author while attached to three submarines and a training facility (Naval Guided Missile School Detachment) during the years 1982-1989.
- [Foley87] Foley, James D., “Interfaces for Advanced Computing”, *Scientific American*, October 1987, pp. 127 ff.

- [Furness91] Furness, Thomas A., III. Course Notes: *Virtual Interface Technology (Virtual Reality)*, Course C3, SIGGRAPH 1991.
- [Gasper91] Gasper, Elan. Demonstration of the product *At Your Service* for the Macintosh on 5 September 1991. Bright Star Technologies, 1450 114th Avenue SE, Suite 200, Bellevue, WA 98004.
- [Gibson84] Gibson, William. *Neuromancer*. New York: Ace Books, 1984.
- [1Gibson86] Gibson, William. *Count Zero*. New York: Ace Science Fiction Books, 1986.
- [2Gibson86] Gibson, William. *Burning Chrome*. London: Grafton Books, 1986.
- [Gibson88] Gibson, William. *Mona Lisa Overdrive*. New York: Bantam Books, 1988.
- [Goodman90] Goodman, Danny. *The Complete HyperCard 2.0 Handbook*. (3rd ed.). Toronto: Bantam Books, 1991.
- [Gore92] Gore, Andrew. "Zeppo to make 3-D marks". *MacWEEK*, Volume 6, Number 10 (9 March 92), p. 1ff.
- [Henderson86] Henderson, D.A. Jr., and Card, S.K. "Rooms: The Use of Multiple Virtual Workspaces To Reduce Space Contention in a Window-Based Graphical User Interface". *ACM Transactions on Graphics*, vol. 5 (3 July 1986), pp. 211 -243.
- [Hollis92] Hollis, Robert. "Voice Technology Beckons". *MacWEEK*, Volume 6, Number 11 (16 March 92), p. 33-4.
- [Kawasaki91] Kawasaki, Guy. *Selling the Dream*. New York: HarperCollins Publishers, 1991.
- [Kelly89] Kelly, Kevin, Heilbrun, Adam, and Stacks, Barbara, "Virtual Reality: An Interview With Jaron Lanier", *Whole Earth Review*, Fall 1989, pp. 108 ff.
- [Krueger91] Krueger, Myron. *Artificial Reality II*. Reading: Addison-Wesley Publishing Company, 1991.
- [Lanier91] Lanier, Jaron, and Nonno, Linda (Co-Chairs), "Applications of Virtual Reality: Reports From The Field", *Computer Graphics*, Vol. 25, Number 4, July 1991, pp.365 ff.
- [Laurel90] Laurel, Brenda, editor, *The Art of Human-Computer Interface Design*. Reading: Addison-Wesley Publishing Company, 1990.

- [Mackinlay91] Mackinlay, Jock D., et. al. "The Perspective Wall: Detail and Context Smoothly Integrated". *Human Factors In Computing Systems* (ACM SIGCHI Conference Proceedings), 1991, pp. 173-9.
- [Mandala91] Vincent, J. V. "Mandala System: Virtual Corridors, Authoring and Virtual Constructions, and Teleconferencing" (Vivid Effects, Inc.), virtual world exhibit demonstrated at ACM SIGGRAPH '91. *Computer Graphics*, Vol. 25, Number 4, July 1991, p. 382.
- [Marsan91] Marsan, Carolyn Duffy. "Navy School Makes Virtual Reality Fly". *Federal Computer Week*, August 12, 1991, p. 8 ff.
- [McManus92] McManus, Neil. "Sketch! back from the drawing board". *MacWEEK* Volume 6 Number 10 (9 March 92), p. 18.
- [Naimark91] Naimark, Michael. "EAT - A Virtual Dining Environment" (Naimark and Company), virtual world exhibit demonstrated at ACM SIGGRAPH '91. *Computer Graphics*, Vol. 25, Number 4, July 1991, p. 382.
- [Nielsen90] Nielsen, Jakob, *Hypertext and Hypermedia*. San Diego: Academic Press, 1990.
- [Rheingold85] Rheingold, Howard, *Tools for Thought: The People and Ideas Behind the NeXT Computer Revolution*. New York, Simon and Schuster, 1985.
- [Rheingold91] Rheingold, Howard, *Virtual Reality*. New York: Summit Books, 1991.
- [Robertson91] Robertson, George G., et. al. "Cone Trees: Animated 3D Visualizations of Hierarchical Information". *Human Factors In Computing Systems* (ACM SIGCHI Conference Proceedings), 1991, pp. 189-194.
- [Schneiderman89] Schneiderman, Ben, and Kearsley, Greg. *Hypertext Hands-On!: An Introduction to a New Way of Organizing and Accessing Information*. Reading: Addison-Wesley Publishing Company, 1989.
- [Stahr91] Stahr, Lisa. "IRIS INDIGO," *IRIS Universe* Number 17, p. 8-11.
- [Teller91] Teller, Seth J., Séquin, Carlo H. "Visibility Preprocessing For Interactive Walkthroughs", *Computer Graphics*, Vol. 25, Number 4, July 1991, pp. 61 ff.
- [Woolf90] Woolf, Beverly, *Intelligent Tutoring Systems: A Survey*. University of Massachusetts (Department of Computer and Information Science), Amherst, Massachusetts (February 1990).

- [Zyda91] Zyda, Michael and Pratt, David, "NPSNET: A 3D Visual Simulator for Virtual World Exploration and Experimentation", *1991 SID International Symposium Digest of Technical Papers*, Volume XXII, 8 May 1991, pp. 361-364.
- [Zyda92] Zyda, Michael J., Pratt, David R., Monahan, James G., Wilson, Kalin P., "NPSNET: Constructing A 3D Virtual World". Paper to appear in the *Proceedings of the 1992 Symposium on Interactive 3D Graphics*, Cambridge, Massachusetts, 29 March 1992.

ANNOTATED BIBLIOGRAPHY

Atwood, Donald J., *DoD Modeling and Simulation Management Plan* (Memorandum from the U.S. Deputy Secretary of Defense), ser. 38917, 21 June 1991. This plan details the objective of DOD with respect to simulation systems, enunciating the need to validate system design using simulation before building more expensive prototypes.

Badler, Norman I., Barsky, Brian, and Zeltzer, David, eds. "Task-level Graphical Simulation: Abstraction, Representation, and Control" by David Zeltzer. *Making Them Move*. (Chapter 1). San Mateo, California: Morgan Kaufmann Publishers, Inc., 1991, pp. 3 ff. Fine articles throughout; Zeltzer's article in Chapter 1 points out the complexities involved in the animation task.

Brand, Stewart, *The Media Lab*. New York: Penguin Books, 1988. This book is brain food, and should be required reading for anyone with an interest in modern computer technology, and even moreso for professionals in the field. It's a required text for my Interactive Computing Systems course.

Bush, Vannevar, "As We May Think". *Atlantic Monthly*, July 1945, pp. 101-108. The original manifesto that lit the hypertext candle; discussion of the Memex system, which resembles a hypertext system on paper.

Card, Stuart K., et. al. "The Information Visualizer: An Information Workspace". *Human Factors In Computing Systems* (ACM SIGCHI Conference Proceedings), 1991, pp. 181-8. The only other system I could find that attempted to merge hypertext and a virtual world.

Gibson, William. *Neuromancer*. New York: Ace Books, 1984. Gives great insights on some of the possibilities for virtual world issues. See also *Burning Chrome* by the same author.

Henderson, D.A. Jr., and Card, S.K. "Rooms: The Use of Multiple Virtual Workspaces To Reduce Space Contention in a Window-Based Graphical User Interface". *ACM Transactions on Graphics*, vol. 5 (3 July 1986), pp. 211 -243. Rooms is the predecessor to 3D/Rooms, part of the Information Visualizer project.

Kawasaki, Guy. *Selling the Dream*. New York: HarperCollins Publishers, 1991. The bible of evangelism for a cause; sheds light on how to convince people that virtual reality and hypertext offer great things for the future, for instance. If I have been successfully persuasive here, this book deserves a lot of the credit. If I have failed, it's not the book's fault.

Kelly, Kevin, Heilbrun, Adam, and Stacks, Barbara, "Virtual Reality: An Interview With Jaron Lanier", *Whole Earth Review*, Fall 1989, pp. 108 ff. Good introduction on body suits, data gloves, and goggles; also has an interesting depiction of virtual world authoring on a Macintosh.

Krueger, Myron. *Artificial Reality II*. Reading: Addison-Wesley Publishing Company, 1991.

Some interesting virtual world concepts; a different perspective than most.

Laurel, Brenda, editor, *The Art of Human-Computer Interface Design*. Reading: Addison-Wesley Publishing Company, 1990.

Mackinlay, Jock D., et. al. "The Perspective Wall: Detail and Context Smoothly Integrated". *Human Factors In Computing Systems* (ACM SIGCHI Conference Proceedings), 1991, pp. 173-9.

Part of the Information Visualizer project reference above with Stuart Card of Xerox PARC. Fascinating and well-written.

Marsan, Carolyn Duffy. "Navy School Makes Virtual Reality Fly". *Federal Computer Week*, August 12, 1991, p. 8 ff.

Nielsen, Jakob, *Hypertext and Hypermedia*. San Diego: Academic Press, 1990. The best introduction to hypertext in the world. Informative and well-written.

Rheingold, Howard, *Tools for Thought: The People and Ideas Behind the NeXT Computer Revolution*. New York, Simon and Schuster, 1985. Good depiction of the NeXT cause and its objectives. More applicable to the multimedia side of the subject than virtual worlds.

Rheingold, Howard, *Virtual Reality*. New York: Summit Books, 1991. A great overview of virtual worlds by a non-technical observer. Brain food.

Robertson, George G., et. al. "Cone Trees: Animated 3D Visualizations of Hierarchical Information". *Human Factors In Computing Systems* (ACM SIGCHI Conference Proceedings), 1991, pp. 189-194. Xerox PARC Information Visualizer project; worked with Card and MacKinlay.

Schneiderman, Ben, and Kearsley, Greg. *Hypertext Hands-On!: An Introduction to a New Way of Organizing and Accessing Information*. Reading: Addison-Wesley Publishing Company, 1989. Another good introduction to hypertext, this book comes with a hypertext version of itself; unfortunately, it's only available for DOS machines.

Stahr, Lisa. "IRIS INDIGO," *IRIS Universe* Number 17, p. 8-11. This short article details the most affordable workstation capable of real-time interactive 3D. The Indigo could be the spark that lights the virtual reality + hypertext revolution.

Woolf, Beverly, *Intelligent Tutoring Systems: A Survey*. University of Massachusetts (Department of Computer and Information Science), Amherst, Massachusetts (February 1990). An extremely insightful look at how ITS's can help the education process.

Zyda, Michael and Pratt, David, "NPSNET: A 3D Visual Simulator for Virtual World Exploration and Experimentation", *1991 SID International Symposium Digest of Technical Papers*, Volume XXII, 8 May 1991, pp. 361-364. Description of NPSNET in detail, and virtual worlds issues.

Zyda, Michael J., Pratt, David R., Monahan, James G., Wilson, Kalin P., "NPSNET: Constructing A 3D Virtual World". Paper to appear in the *Proceedings of the 1992 Symposium on Interactive 3D Graphics*, Cambridge, Massachusetts, 29 March 1992. Current state of NPSNET.

INITIAL DISTRIBUTION LIST

Defense Technical Information Center Cameron Station Alexandria, VA 22304-6145	2
Dudley Knox Library Code 52 Naval Postgraduate School Monterey, CA 93943	2
Dr. Robert McGhee, Chairman, Code CSMz Computer Science Department Naval Postgraduate School Monterey, CA 93943	1
Dr. Michael Zyda, Code CS/Zk Computer Science Department Naval Postgraduate School Monterey, CA 93943	2
Dr. Yuh-jeng Lee, Code CS/Le Computer Science Department Naval Postgraduate School Monterey, CA 93943	1
Mr. David R. Pratt, Code CS/Pr Computer Science Department Naval Postgraduate School Monterey, CA 93943	1
LCDR John A. Daley, Code CS/Da Computer Science Department Naval Postgraduate School Monterey, CA 93943	3
Mr. Steven E. Anderson c/o NSWC Dahlgren Division Systems Accuracy Branch (G12) Dahlgren, VA 22448-5000	1

LCDR Don Brutzman, Code OR/Br 1
Computer Science Department
Naval Postgraduate School
Monterey, CA 93943

CAPT Jack Jensen, Code 00/FNOC 1
Fleet Numerical Oceanographic Center
Monterey, CA 93943

Chuck Lombardo, Computer Systems Programmer 1
Computer Science Department
Naval Postgraduate School
Monterey, CA 93943

Thesis
D1418 Daley
c.1 Hyper-NPSNET.



DUDLEY KNOX LIBRARY



3 2768 00034210 9